

A Hierarchical Model for Simultaneous Detection and Estimation in Multi-subject fMRI Studies

David Degras^a and Martin A. Lindquist^{b*}

^a*Department of Mathematical Sciences, DePaul University, USA.*

^b*Department of Biostatistics, Johns Hopkins University, USA.*

April 1, 2014

*Corresponding author: 615 N. Wolfe Street, E3634; Baltimore, MD 21205, e-mail: mlindqui@jhsph.edu.
This research was supported by NIH grant R01EB016061.

Abstract

In this paper we introduce a new hierarchical model for the simultaneous detection of brain activation and estimation of the shape of the hemodynamic response in multi-subject fMRI studies. The proposed approach circumvents a major stumbling block in standard multi-subject fMRI data analysis, in that it both allows the shape of the hemodynamic response function to vary across region and subjects, while still providing a straightforward way to estimate population-level activation. An efficient estimation algorithm is presented, as is an inferential framework that not only allows for tests of activation, but also for tests for deviations from some canonical shape. The model is validated through simulations and application to a multi-subject fMRI study of thermal pain.

1 INTRODUCTION

Depending on their scientific goals, researchers in functional magnetic resonance imaging (fMRI) often choose modeling strategies with the intent to either *detect* the magnitude of activation in a certain brain region, or *estimate* the shape of the hemodynamic response associated with the task being performed [Poldrack et al., 2011]. While most of the focus in neuroimaging to date has been on detection [Lindquist, 2008], the magnitude of evoked activation cannot be accurately measured without either assuming or measuring timing and shape information as well. In practice, many statistical models of fMRI data attempt to simultaneously incorporate information about the shape, timing, and magnitude of task-evoked hemodynamic responses.

As an example, consider the general linear model (GLM) approach [Worsley and Friston, 1995], which is arguably the dominant approach towards analyzing fMRI data. It models the fMRI time series as a linear combination of several different signal components and tests whether activity in a brain region is related to any of them. Typically the shape of the hemodynamic response is assumed *a priori*, using a canonical hemodynamic response function (HRF) [Friston et al., 1998, Glover, 1999], and the focus of the analysis is on obtaining the magnitude of the response across the brain. However, it is well-known that the shape of the HRF varies both across space and subjects [Aguirre et al., 1998, Schacter et al., 1997, Handwerker et al., 2004, Badillo et al., 2013a]; thus assuming a constant shape across all voxels and subjects may give rise to significant bias in large parts of the brain [Lindquist and Wager, 2007, Lindquist et al., 2009]. The constant HRF assumption can be relaxed by expressing it as a linear combination of several known basis functions. This can be done within the GLM framework by convolving the same stimulus function with multiple canonical waveforms and including them as multiple columns of the design matrix for each condition. The coefficients for an event type constructed using different basis functions can then be combined to fit the evoked HRF in that particular area of the brain.

The ability to use basis sets to capture variations in hemodynamic responses depends both on the number and shape of the reference waveforms that are used in the model. For example, the finite impulse response (FIR) basis set, consists of one free parameter for every time-point following stimulation in every cognitive event-type that is modeled [Glover, 1999, Goutte et al., 2000]. Therefore, it can be used to estimate HRFs of arbitrary shape for each event type in every voxel of the brain. Another, perhaps more common approach, is to use the canonical HRF together with its temporal and dispersion derivatives to allow for small shifts in both the onset and width of the HRF. Other choices of basis sets include those composed of principal components [Aguirre et al., 1998, Woolrich et al., 2004], cosine functions [Zarahn, 2002], radial basis functions [Riera et al., 2004], spectral basis sets [Liao et al., 2002] and inverse logit functions [Lindquist and Wager, 2007]. For a critical evaluation of a number of commonly used basis sets, see [Lindquist and Wager, 2007] and [Lindquist et al., 2009].

Though basis sets allow the constant HRF assumption to be relaxed in the GLM framework, they are still not without problems, particularly when performing multi-subject analysis. Most analyses of multi-subject fMRI data involve two separate models. A first-level GLM

is performed separately on each subject’s data, providing subject-specific contrasts of parameter estimates. A second-level model is thereafter used to provide population inference on whether the contrasts are significantly different from zero and assess the effects of any group-level predictors, such as group status or behavioral performance. However, it is problematic to define appropriate first-level contrasts that truly capture the behavior we are interested in detecting, when multiple basis sets are included in the model. For example, when each condition consists of multiple basis functions it is not self-evident how to properly define a relevant contrast comparing the difference in activation between the two conditions.

There have been a number of suggestions about how to deal with this issue, most notably using only the “main” basis set [Kiehl and Liddle, 2001]. In the case of the canonical HRF and its derivatives, this entails only using the coefficient corresponding to the canonical HRF, and treating the coefficients corresponding to the derivatives as nuisance parameters. While this works relatively well for small deviations from the HRFs canonical form, it quickly falls apart as the shape begins to differ. To circumvent this problem [Calhoun et al., 2004] suggested using the norm of the coefficients for the canonical HRF and its derivatives. Another suggestion [Lindquist et al., 2009] is to re-create the HRF after estimation and use the resulting amplitude as the contrast of interest.

In this work we introduce a new approach towards multi-subject analysis of fMRI data that enables us to simultaneously estimate the specific shape of the HRF for a subject in a given voxel, and to obtain a population-level estimate of the magnitude of activation. This approach offers the flexibility of basis sets while retaining the simplicity of multi-subject inference with a single canonical HRF. We also provide an inferential framework that allows us to test both for activations as well as for any differences in HRF shape from some canonical form.

The idea of performing joint estimation and detection is not new in the neuroimaging literature. For example, Makni and colleagues [Makni et al., 2005, Makni et al., 2008] have suggested a Bayesian approach towards the detection of brain activity that uses a mixture of two Gaussian distributions as a prior on a latent neural response, whereas the hemodynamic impulse response is constrained to be smooth using a Gaussian prior. In this model all parameters of interest are estimated from the posterior distribution using Gibbs sampling. Later work has provided a number of interesting extensions of this model, including to spatial mixture models [Vincent et al., 2010] and reformulating it as a missing data problem that allows for a simplified estimation procedure [Chaari et al., 2013].

Our suggested model takes a different approach. Here the HRF is modeled as a linear combination of B-spline functions (see e.g., [Genovese, 2000] for an early use of spline functions in model fitting). We assume that subject-level HRFs are random draws from a population-level distribution and that for any given voxel the population average response across stimuli will vary only in scale. We provide an efficient algorithm for estimating the model parameters as well as inferential methods. The latter includes both tests of activation and for deviations in the HRF from some canonical form.

This paper is organized as follows. In Section 2 we introduce our hierarchical model. In Sections 3 and 4 we outline an efficient algorithm for estimating the model parameters and

performing inference on them. In Sections 5 and 6 we evaluate the performance of the model on a series of simulated data sets and data from an experiment studying the effects of thermal pain. Both the simulated and real data were previously used in a large study of flexible HRF modeling procedures [Lindquist et al., 2009]. The proposed method is shown to outperform each of the previously tested approaches, which include the canonical GLM plus its derivatives, the smooth FIR model, and the inverse logit model. We conclude with a discussion of the suggested approach.

2 METHODS

2.1 The Hierarchical Model

In this section we outline the proposed hierarchical model for simultaneous estimation and detection. For simplicity, we assume only one session per subject and a single experimental group. To study multiple groups, it suffices to apply the model below and the estimation procedure of section 2.2 separately to each group. Similarly, the inference method of section 2.3 easily extends to multiple samples. We assume that all scans have been acquired at the same repetition time Δ and registered to a standard stereotactic space. For the j^{th} subject ($1 \leq j \leq n$), we model the BOLD response at the v^{th} voxel ($1 \leq v \leq V$) and t^{th} scan ($1 \leq t \leq T_j$) as

$$\underbrace{y_j(v, t)}_{\text{BOLD}} = \sum_{l=1}^L \int_0^{t\Delta} \underbrace{s_{jl}(t\Delta - \tau)}_{\text{stimulus}} \underbrace{h_{jl}(v, \tau)}_{\text{HRF}} d\tau + \sum_{\nu=1}^q d_{j\nu}(v) \underbrace{\varphi_{j\nu}(t)}_{\text{nuisance}} + \underbrace{\varepsilon_j(v, t)}_{\text{noise}}. \quad (1)$$

Here the response consists of a linear combination of stimulus-induced signal of interest (represented by the first sum on the right-hand side of the equation), nuisance signal (the second sum) and noise. The stimulus function s_{jl} is a stick-function that has baseline at zero and takes the value one whenever stimuli of the l^{th} type are presented. The nuisance signals $\varphi_{j\nu}$, $1 \leq \nu \leq q$, typically include scanner drift, represented by polynomial and/or cosine basis sets, and physiological noise such as head motion, heart beat, and respiration. The corresponding nuisance parameters $d_{j\nu}$ must be estimated from the data. The subject-level HRF h_{jl} decomposes into

$$h_{jl}(v, \tau) = h_l(v, \tau) + \eta_{jl}(v, \tau), \quad (2a)$$

where h_l is a population-level HRF (fixed effect) and η_{jl} is a subject-specific (random) effect. Hence, each subject-level HRF is assumed to be a random draw from a population with mean h_l . Thus, similar to other recent methods (e.g. [Sanyal and Ferreira, 2012] and [Zhang et al., 2013]) this allows us to borrow strength across subjects to improve subject-specific estimation. The functions $h_l(v, \cdot)$ and $\eta_{jl}(v, \cdot)$ are represented using a set of basis functions $(B_k)_{1 \leq k \leq K}$:

$$h_l(v, \tau) = \sum_{k=1}^K \gamma_{lk}(v) B_k(\tau), \quad \eta_{jl}(v, \tau) = \sum_{k=1}^K \xi_{jlk}(v) B_k(\tau), \quad (2b)$$

so estimation reduces to solving for the coefficients $\gamma_{lk}(v)$ and $\xi_{jlk}(v)$. In practice we suggest specifying the B_k as B-splines with regularly spaced knots over a time interval where the HRF is believed to be non-zero, say in the range between 0 and 30 seconds. B-spline basis sets have several desirable features. First, the coefficients of a function in a B-spline basis are very close to the function itself, i.e., the function values at the knots (possibly up to a scaling factor). For this reason, B-spline coefficients are immediately interpretable and inference of local features of the HRF is greatly facilitated. In addition, the compact support of B-splines typically induces sparsity in the design matrices and thus reduces the computational load.

Noting that the shape of the HRF at a given brain location is mostly determined by physiological factors that are independent of the nature of the stimulus, we further assume that the population-level HRFs $h_l(v, \cdot)$, $1 \leq l \leq L$, have the same shape for all conditions and differ only in scale:

$$\gamma_{lk}(v) = \beta_l(v) \gamma_k(v). \quad (2c)$$

Here the terms $\gamma_k(v)$, $1 \leq k \leq K$, determine the shape of the HRF, while $\beta_l(v)$ determines the amplitude of the response to stimulus of the l^{th} type. To make these parameters identifiable, we impose the scale constraint $\sum_k \gamma_k(v)^2 = 1$ and the orientation constraint $\gamma_{k_0}(v) > 0$, with $k_0 = \arg \max_k |\gamma_k(v)|$. Assumption (2c) considerably reduces the number of HRF parameters from KL to $K + L$ (the same modeling assumption is used in e.g., [Makni et al., 2005]), which allows one to use a reasonably large number of basis functions while maintaining a good estimation accuracy. On the other hand, the product form of (2c) makes model (1) nonlinear with respect to the parameters $\beta_l(v)$ and $\gamma_k(v)$. Note that the shape coefficients $\gamma_k(v)$ are well defined only if at least one of the HRFs $h_l(v, \cdot)$ is not identically zero, that is, if at least one experimental condition induces an activation at voxel v .

The random effects $\xi_{jlk}(v)$ in (2b) represent subject-specific deviations from the group-level HRF coefficients $\gamma_{lk}(v)$. They are assumed to be independent across subjects and conditions. In addition they are assumed to be Gaussian random vectors with mean zero and correlation structure that is stationary in time and constant in space:

$$\text{Cov}(\xi_{jlk}(v), \xi_{j'l'k'}(v)) = (\delta_{jj'} \delta_{ll'}) \sigma_{\xi l}^2(v) \rho_{\xi l}(|k - k'|). \quad (3)$$

In the previous equation, $\sigma_{\xi l}^2(v)$ denotes the common variance of the $\xi_{jlk}(v)$ ($1 \leq j \leq n$, $1 \leq k \leq K$), $\rho_{\xi l}$ is an autocorrelation function, and δ is the Kronecker delta ($\delta_{xy} = 1$ if $x = y$, $\delta_{xy} = 0$ if $x \neq y$). Note that $\rho_{\xi l}(0) = 1$.

To specify the dependence structure of the noise component ε , we consider a partition of the spatial domain \mathcal{D} into neuro-anatomic parcels $\mathcal{D}_1, \dots, \mathcal{D}_M$ (e.g., Brodmann areas or any suitable brain atlas). We assume that for each voxel v of a parcel \mathcal{D}_m , $\varepsilon_j(v, \cdot)$ is a stationary Gaussian AR(p) process whose variance $\sigma_{\varepsilon m}^2$ and structural parameters $\boldsymbol{\theta}_{\varepsilon m} = (\theta_{\varepsilon 1m}, \dots, \theta_{\varepsilon pm})$ only depend on m . In other words, $\varepsilon_j(v, t) = \sum_{k=1}^p \theta_{km} \varepsilon_j(v, t-k) + e_j(v, t)$, where the $e_j(v, t)$, $1 \leq t \leq N_j$, are i.i.d. normal innovations. The specification of the noise dependence at the parcel level reflects the belief that this noise is spatially smooth. An alternative way to characterize the spatial smoothness of ε would be to model the AR parameters σ_{ε}^2 and $\boldsymbol{\theta}_{\varepsilon}$ as smooth functions of v .

Model Summary

For the j^{th} subject and v^{th} voxel, the BOLD time course $\mathbf{y}_j(v) = (y_j(v, 1), \dots, y_j(v, T_j))'$ can be expressed in matrix form as

$$\mathbf{y}_j(v) = \mathbf{X}_j(\boldsymbol{\beta}(v) \otimes \boldsymbol{\gamma}(v)) + \mathbf{X}_j\boldsymbol{\xi}_j(v) + \boldsymbol{\Phi}_j\mathbf{d}_j(v) + \boldsymbol{\varepsilon}_j(v). \quad (4)$$

As in the standard GLM, the design matrix \mathbf{X}_j is the convolution of the stimulus functions s_{jl} with the basis functions B_k . To be precise, $\mathbf{X}_j = (\mathbf{X}_{j1}, \dots, \mathbf{X}_{jL})$ with $\mathbf{X}_{jl} = (\mathbf{x}_{j1l}, \dots, \mathbf{x}_{jKl})$ and $\mathbf{x}_{jkl} = (\int_0^\Delta B_k(\tau)s_{jl}(\Delta - \tau)d\tau, \dots, \int_0^{T_j\Delta} B_k(\tau)s_{jl}(T_j\Delta - \tau)d\tau)'$. We have also written $\boldsymbol{\beta}(v) = (\beta_1(v), \dots, \beta_L(v))'$ and $\boldsymbol{\gamma}(v) = (\gamma_1(v), \dots, \gamma_K(v))'$ for the amplitude and scale parameters of the population-level HRF; $\boldsymbol{\xi}_j(v) = (\xi_{j1}(v), \dots, \xi_{jL}(v))'$ and $\boldsymbol{\xi}_{jl}(v) = (\xi_{j1l}(v), \dots, \xi_{jKl}(v))'$ for the subject-specific effects; $\boldsymbol{\Phi}_j = (\varphi_{j\nu}(t))_{1 \leq \nu \leq q, 1 \leq t \leq T_j}$ and $\mathbf{d}_j(v) = (d_{j1}(v), \dots, d_{jq}(v))'$ for the nuisance signals; and $\boldsymbol{\varepsilon}_j(v) = (\varepsilon_j(v, 1), \dots, \varepsilon_j(v, T_j))'$ for the noise. The symbol \otimes denotes the Kronecker product.

The vector $\mathbf{y}_j(v)$ has a multivariate normal distribution with mean and covariance

$$\begin{aligned} \boldsymbol{\mu}_j(v) &= \mathbf{X}_j(\boldsymbol{\beta}(v) \otimes \boldsymbol{\gamma}(v)) + \boldsymbol{\Phi}_j\mathbf{d}_j(v), \\ \mathbf{V}_j(v) &= \mathbf{X}_j(\mathbf{D}_\xi(v) \otimes \mathbf{I}_K)\mathbf{T}_\xi\mathbf{X}_j' + \mathbf{V}_{\varepsilon jm}, \end{aligned} \quad (5)$$

where $\mathbf{D}_\xi(v) = \text{diag}(\sigma_{\xi 1}^2(v), \dots, \sigma_{\xi L}^2(v))$, $\mathbf{T}_\xi = \text{diag}(\mathbf{T}_{\xi 1}, \dots, \mathbf{T}_{\xi L})$ is a block diagonal matrix with $\mathbf{T}_{\xi l} = (\rho_{\xi l}(|k - k'|))_{1 \leq k, k' \leq K}$, \mathbf{I}_K is the $K \times K$ identity matrix, and $\mathbf{V}_{\varepsilon jm}$ is the covariance matrix of $\boldsymbol{\varepsilon}_j(v)$ for $v \in \mathcal{D}_m$. Here the first term of $\mathbf{V}_j(v)$ corresponds to between-subject variation, while the second term is the within-subject variation.

Table 1 summarizes all model notations.

2.2 Estimation

In this section we outline our procedure for estimating the parameters of our model. Our main objective is to estimate the population HRF at each voxel while taking into account the covariance structure of the data. We formulate this objective mathematically as a generalized, penalized, and constrained least squares problem. Since the corresponding objective function is non-convex, our procedure is only guaranteed to yield a local minimum. It is therefore important to select good starting values for the procedure, which we do by constructing a consistent pilot estimator of the HRF. We then provide consistent estimators of the data covariance parameters, after which we optimize the objective function to obtain the final HRF estimates. The entire procedure is performed in the following five steps, each described in detail in a subsequent subsection:

1. For each voxel v , estimate the HRF parameters $\gamma_{kl}(v)$ by Penalized Least Squares (PLS). This pilot estimation does not exploit the HRF shape assumption (2c) and does not account for the covariance structure of the data.
2. For each parcel \mathcal{D}_m , estimate the parameters $\sigma_{\varepsilon m}^2$ and $\boldsymbol{\theta}_{\varepsilon m}$ of the AR noise process ε by solving the Yule-Walker equations associated with the predicted errors $\hat{\varepsilon}_j(v)$. The $\hat{\varepsilon}_j(v)$ are obtained from a least squares fit on the residuals of step 1.

Table 1: Notations

j	Subject index
k	Basis function index
l	Condition index
v	Voxel index
n	Sample size
K	Number of basis functions
L	Number of experimental conditions
V	Number of voxels
B_k	Basis function
T_j	Number of brain scans for subject j
$\mathbf{y}_j(v)$	fMRI time course for subject j at voxel v
\mathbf{X}_j	Design matrix for subject j
\mathbf{X}_{jl}	Design matrix for subject j , condition l
$\beta(v)$	Amplitude parameters for the population HRF at voxel v
$\gamma(v)$	Shape parameters for the population HRF at voxel v
$\xi_j(v)$	Deviation of subject j from population HRF at voxel v
$\xi_{jl}(v)$	Deviation of subject j from population HRF for condition l and voxel v
Φ_j	Matrix of nuisance signals for subject j
$\mathbf{d}_j(v)$	nuisance coefficients for subject j and voxel v
$\epsilon_j(v)$	Noise vector for subject j at voxel v
$\mathbf{D}_\xi(v)$	Variance coefficients of subject effects at voxel v
\mathbf{T}_ξ	Correlation matrix for subject effects at voxel v
$\mathbf{T}_{\xi l}$	Correlation matrix for subject effects for condition l and voxel v
$\mathbf{V}_{\epsilon jm}$	Covariance matrix of the noise for subject j and parcel m
$\mathbf{V}_j(v)$	Covariance matrix of $\mathbf{y}_j(v)$

3. Estimate the temporal correlation parameters $\rho_{\xi l}(k)$ of the subject random effects by Maximum Likelihood (ML). The ML estimates are obtained separately on a small sample of voxels and pooled with a suitable statistic (e.g., trimmed mean or median).
4. For each voxel v , estimate the between-subject variance $\sigma_\xi^2(v)$ by Variance Least Squares (VLS).
5. For each voxel v , estimate $\beta(v)$ and $\gamma(v)$ again using a generalized, penalized, and constrained least squares approach.

Step 1: Pilot estimation of the HRF

For each voxel v , the HRF scale and shape coefficients $\beta(v)$ and $\gamma(v)$ are first estimated by penalized least squares (PLS):

$$\min_{\mathbf{h}, \mathbf{d}} \left\{ \sum_{j=1}^n \|\mathbf{y}_j(v) - \mathbf{X}_j \mathbf{h} - \Phi_j \mathbf{d}_j\|^2 + n\lambda_0 \mathbf{h}'(\mathbf{I}_L \otimes \mathbf{P})\mathbf{h} \right\}, \quad (6)$$

where \mathbf{h} is a vector of length KL that estimates the HRF coefficients $\gamma_{lk}(v)$ and $\mathbf{d} = (\mathbf{d}'_1, \dots, \mathbf{d}'_n)'$ is a vector of length nq that estimates the nuisance signals. The matrix \mathbf{P} penalizes departures of \mathbf{h} from a linear space of “reasonable” HRFs (e.g., the canonical HRF and its temporal derivative). More precisely, let Ψ be a matrix whose columns contain the coefficients of a few realistic HRFs in the function basis $(B_k)_{1 \leq k \leq K}$. Then $\mathbf{P} = \mathbf{I}_K - \Psi(\Psi'\Psi)^{-1}\Psi'$ is the projection on the orthogonal space of Ψ . The smoothing parameter $\lambda_0 > 0$ determines the tradeoff between fitting the data and closeness to Ψ . It can be selected manually or, for example, by k -fold cross-validation with the subjects randomly partitioned in k subsamples.

Note that the minimization problem (6) is unconstrained and does not rely on (2c). Its solutions are $\hat{\mathbf{h}}(v) = (\sum_{j=1}^n \mathbf{X}'_j \mathbf{R}_j \mathbf{X}_j + n\lambda_0(\mathbf{I}_L \otimes \mathbf{P}))^{-1} \sum_{j=1}^n \mathbf{X}'_j \mathbf{R}_j \mathbf{y}_j(v)$ and $\hat{\mathbf{d}}_j(v) = (\Phi'_j \Phi_j)^{-1} \Phi'_j (\mathbf{y}_j(v) - \mathbf{X}_j \hat{\mathbf{h}}(v))$, $1 \leq j \leq n$, where $\mathbf{R}_j = \mathbf{I}_T - \Phi_j(\Phi'_j \Phi_j)^{-1} \Phi'_j$ is the projection matrix on the orthogonal space of Φ_j . The estimator $\hat{\mathbf{h}}(v)$ is consistent and asymptotically unbiased for the $\gamma_{lk}(v)$ as the sample size $n \rightarrow \infty$ and the smoothing parameter $\lambda_0 \rightarrow 0$. In view of (2c), it follows that $\hat{\beta}_0(v) = (\|\hat{\mathbf{h}}_1(v)\|, \dots, \|\hat{\mathbf{h}}_L(v)\|)'$ is a consistent estimator of $\beta(v)$, where $\hat{\mathbf{h}}(v) = (\hat{\mathbf{h}}_1(v)', \dots, \hat{\mathbf{h}}_L(v)')'$ and the vectors $\hat{\mathbf{h}}_l(v)$, $1 \leq l \leq L$, have length K . Similarly, the scaled average $\hat{\gamma}_0(v) = \sum_{l=1}^L \hat{\mathbf{h}}_l(v) / \|\sum_{l=1}^L \hat{\mathbf{h}}_l(v)\|$ consistently estimates $\gamma(v)$ when the latter is well-defined, i.e., when the voxel v is activated by at least one experimental condition.

Step 2: Estimation of the noise structure

We turn to the estimation of the noise parameters $\sigma_{\varepsilon m}^2$ and $\theta_{\varepsilon m}$ in each parcel \mathcal{D}_m , $1 \leq m \leq M$. For each subject $1 \leq j \leq n$ and voxel $v \in \mathcal{D}_m$, consider the residual vector $\mathbf{r}_j(v) = \mathbf{y}_j(v) - \mathbf{X}_j(\hat{\beta}_0(x) \otimes \hat{\gamma}_0(v)) - \Phi_j \hat{\mathbf{d}}(v) = \mathbf{R}_j(\mathbf{y}_j(v) - \mathbf{X}_j(\hat{\beta}_0(x) \otimes \hat{\gamma}_0(v)))$ resulting from step 1. Given (4) and the consistency of $\hat{\mathbf{h}} = \hat{\mathbf{h}}(v)$ in step 1, it holds that $\mathbf{r}_j(v) \approx \mathbf{X}_j \boldsymbol{\xi}_j + \boldsymbol{\varepsilon}_j(v)$ for n large enough and λ_0 small. The random effects $\boldsymbol{\xi}_j(v)$ and $\boldsymbol{\varepsilon}_j(v)$ can thus be predicted by least squares based on $\mathbf{r}_j(v)$, yielding

$$\hat{\boldsymbol{\xi}}_j(v) = (\mathbf{X}'_j \mathbf{X}_j)^{-1} \mathbf{X}'_j \mathbf{r}_j(v), \quad \hat{\boldsymbol{\varepsilon}}_j(v) = \mathbf{r}_j(v) - \mathbf{X}_j \hat{\boldsymbol{\xi}}_j(v).$$

If the design matrix \mathbf{X}_j is not full rank, the inverse $(\mathbf{X}'_j \mathbf{X}_j)^{-1}$ in the above formula is not defined and can be replaced by its pseudoinverse $(\mathbf{X}'_j \mathbf{X}_j)^+$. We then solve the Yule-Walker equations (see e.g., [Brockwell and Davis, 2006, p. 239]) associated with $\hat{\boldsymbol{\varepsilon}}_j(v)$, producing consistent estimates of $\sigma_{\varepsilon m}^2$ and $\theta_{\varepsilon m}$ for each subject $1 \leq j \leq n$ and voxel $v \in \mathcal{D}_m$. By taking the medians of these estimates across subjects and voxels, we obtain robust estimates $\hat{\sigma}_{\varepsilon m}^2$ and $\hat{\theta}_{\varepsilon m}$.

Step 3: Estimation of the temporal dependence in subject-specific effects

We estimate the temporal correlation parameters $\rho_{\ell l}(k)$, $1 \leq k \leq K-1$, $1 \leq l \leq L$, by Maximum Likelihood (ML). For computational efficiency, ML estimates are separately produced at a small number of voxels and aggregated with a suitable statistic such as the median or trimmed mean. In practice, we propose to select a random sample of about 1000

voxels for the ML estimation. Following a common usage, we first perform a few iterations of the EM algorithm to provide good starting values for the optimization of the likelihood function.

In the rest of this section, we fix a voxel and omit ther index v for conciseness. Writing $\boldsymbol{\sigma}_\xi^2 = (\sigma_{\xi 1}^2, \dots, \sigma_{\xi L}^2)'$ and $\boldsymbol{\rho}_\xi = (\rho_{\xi 1}(1), \rho_{\xi 1}(2), \dots, \rho_{\xi L}(K-1))'$, we optimize the log-likelihood function (multiplied by -2)

$$\begin{aligned} L(\boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{d}, \boldsymbol{\sigma}_\xi^2, \boldsymbol{\rho}_\xi, \sigma_{\varepsilon m}^2, \boldsymbol{\theta}_{\varepsilon m}) \\ = \sum_{j=1}^n \ln |\mathbf{V}_j| + \sum_{j=1}^n (\mathbf{y}_j - \mathbf{X}_j(\boldsymbol{\beta} \otimes \boldsymbol{\gamma}) - \boldsymbol{\Phi}_j \mathbf{d}_j)' \mathbf{V}_j^{-1} (\mathbf{y}_j - \mathbf{X}_j(\boldsymbol{\beta} \otimes \boldsymbol{\gamma}) - \boldsymbol{\Phi}_j \mathbf{d}_j) \end{aligned} \quad (7)$$

with respect to $\boldsymbol{\sigma}_\xi^2$ and $\boldsymbol{\rho}_\xi$ while fixing $\boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{d}, \sigma_{\varepsilon m}^2$, and $\boldsymbol{\theta}_{\varepsilon m}$ to their previously estimated values. Note that although we are only concerned here with the estimation of $\boldsymbol{\rho}_\xi$, the likelihood must also be optimized with respect to $\boldsymbol{\sigma}_\xi^2$. These variance parameters, which must be estimated at each voxel, will be assessed more efficiently in step 4. The implementation of the EM algorithm and likelihood optimization is described in Appendix A. More details on the EM algorithm for linear mixed models can be found in e.g., ([Pawitan, 2001], chap. 12).

Step 4: Estimation of the between-subjects variance

For each voxel v , we estimate the between-subjects variances $\boldsymbol{\sigma}_\xi^2(v)$ by a variance least squares approach (e.g., [Demidenko, 2004], chap. 3) that consists of minimizing the distance between the residual covariance matrix and the theoretical covariance matrix:

$$\min_{\boldsymbol{\sigma}_\xi^2} \sum_j \left\| \mathbf{r}_j(v) \mathbf{r}_j'(v) - \mathbf{X}_j(\mathbf{D}_\xi(v) \otimes \mathbf{I}_K) \hat{\mathbf{T}}_\xi \mathbf{X}_j' - \hat{\mathbf{V}}_{\varepsilon jm} \right\|_F^2 \quad (8)$$

subject to the constraint $\sigma_{\xi l}^2 \geq 0$ for $1 \leq l \leq L$. Recall that the residuals $\mathbf{r}_j(v)$ are defined in step 2 of this section, $\mathbf{D}_\xi = \text{diag}(\sigma_{\xi 1}^2, \dots, \sigma_{\xi L}^2)$, and $\hat{\mathbf{V}}_{\varepsilon jm}$ and $\hat{\mathbf{T}}_\xi$ are the estimates of $\mathbf{V}_{\varepsilon jm}$ and \mathbf{T}_ξ obtained in steps 2-3. The notation $\|\mathbf{A}\|_F$ stands for the Frobenius norm $\text{tr}(\mathbf{A}'\mathbf{A})^{1/2}$ of a matrix \mathbf{A} . Problem (8) is a standard quadratic programming problem that expressed more simply as

$$\min_{\boldsymbol{\sigma}_\xi^2} \left\{ (\boldsymbol{\sigma}_\xi^2)' \mathbf{A} \boldsymbol{\sigma}_\xi^2 - 2\mathbf{b}' \boldsymbol{\sigma}_\xi^2 \right\}, \quad (9)$$

where \mathbf{A} is a $L \times L$ matrix with (l, l') entry $\sum_{j=1}^n \text{tr}(\mathbf{X}_{jl} \hat{\mathbf{T}}_{\xi l} \mathbf{X}_{jl}' \mathbf{X}_{jl'} \hat{\mathbf{T}}_{\xi l'} \mathbf{X}_{jl'}')$ and \mathbf{b} is a vector of length L with l^{th} entry $\sum_{j=1}^n \{ \mathbf{r}_j'(v) \mathbf{X}_{jl} \hat{\mathbf{T}}_{\xi l} \mathbf{X}_{jl}' \mathbf{r}_j(v) - \text{tr}(\hat{\mathbf{T}}_{\xi l} \mathbf{X}_{jl}' \hat{\mathbf{V}}_{\varepsilon jm} \mathbf{X}_{jl}) \}$. The solutions to (8)-(9) can be computed by various methods (e.g., [Nocedal and Wright, 2006], p. 449) that are widely available in software packages.

Step 5: Generalized least squares estimation of the HRF

The pilot estimation of the HRF can be improved upon in two ways: (i) by accounting for the dependence structure of the BOLD signal, and (ii) by imposing the form (2c)

to the HRF estimates. To integrate these features in the estimation, we use a penalized, constrained, generalized least squares approach. For a given voxel v , let $\hat{\mathbf{V}}_j(v) = \mathbf{X}_j(\hat{\mathbf{D}}_\xi(v) \otimes \mathbf{I}_K)\hat{\mathbf{T}}_\xi\mathbf{X}_j' + \hat{\mathbf{V}}_{\varepsilon jm}$ be the estimate of $\mathbf{V}_j(v)$ resulting from steps 1-4. We seek to solve

$$\min_{\boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{d}} \left\{ \sum_{j=1}^n \left\| \mathbf{y}_j(v) - \mathbf{X}_j(\boldsymbol{\beta} \otimes \boldsymbol{\gamma}) - \boldsymbol{\Phi}_j \mathbf{d}_j \right\|_{\hat{\mathbf{V}}_j^{-1}(v)}^2 + n\lambda \boldsymbol{\gamma}' \mathbf{P} \boldsymbol{\gamma} \right\} \quad (10)$$

under the constraint $\|\boldsymbol{\gamma}\|^2 = 1$. Like in the pilot estimation, the nuisance parameter \mathbf{d} can be eliminated from (10). To that intent, let $\tilde{\mathbf{R}}_j(v) = \mathbf{I}_T - \boldsymbol{\Phi}_j(\boldsymbol{\Phi}_j' \hat{\mathbf{V}}_j^{-1}(v) \boldsymbol{\Phi}_j)^{-1} \boldsymbol{\Phi}_j' \hat{\mathbf{V}}_j^{-1}(v)$ be the projection matrix on the orthogonal space of $\boldsymbol{\Phi}_j$ in the metric $\hat{\mathbf{V}}_j^{-1}(v)$. Then (10) is equivalent to $\min_{\boldsymbol{\beta}, \boldsymbol{\gamma}} \left\{ \sum_{j=1}^n \left\| \tilde{\mathbf{R}}_j(v) \mathbf{y}_j(v) - \tilde{\mathbf{R}}_j(v) \mathbf{X}_j(\boldsymbol{\beta} \otimes \boldsymbol{\gamma}) \right\|_{\hat{\mathbf{V}}_j^{-1}(v)}^2 + n\lambda \boldsymbol{\gamma}' \mathbf{P} \boldsymbol{\gamma} \right\}$.

Because of the tensor product $\boldsymbol{\beta} \otimes \boldsymbol{\gamma}$ and the quadratic constraint $\|\boldsymbol{\gamma}\|^2 = 1$, problem (10) is nonlinear and has no closed-form solutions. However, (10) is a separable least squares problem: for a fixed $\boldsymbol{\gamma}$, solving (10) with respect to $\boldsymbol{\beta}$ reduces to a generalized least squares problem that admits a closed-form solution. For a fixed $\boldsymbol{\beta}$, solving (10) with respect to $\boldsymbol{\gamma}$ is a quadratically constrained quadratic program that requires little more than a singular value decomposition. As a result, (10) can be efficiently solved in an iterative way.

For conciseness, we omit the index v from notations in the remainder of the section. Let $\mathbf{M} = \sum_{j=1}^n \mathbf{X}_j' \tilde{\mathbf{R}}_j' \hat{\mathbf{V}}_j^{-1} \tilde{\mathbf{R}}_j \mathbf{X}_j$ and $\boldsymbol{\eta} = \sum_{j=1}^n \mathbf{X}_j' \hat{\mathbf{V}}_j^{-1} \tilde{\mathbf{R}}_j \mathbf{y}_j$. The solutions $\hat{\boldsymbol{\beta}}$ and $\hat{\boldsymbol{\gamma}}$ of (10) are obtained by cycling through the following equations until convergence:

$$\hat{\boldsymbol{\beta}} = [(\mathbf{I}_L \otimes \hat{\boldsymbol{\gamma}})' \mathbf{M} (\mathbf{I}_L \otimes \hat{\boldsymbol{\gamma}})]^{-1} (\mathbf{I}_L \otimes \hat{\boldsymbol{\gamma}})' \boldsymbol{\eta}, \quad (11)$$

$$\hat{C} = \arg \min_C \left\{ \boldsymbol{\eta}' (\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K) \left[(\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K)' \mathbf{M} (\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K) + n\lambda \mathbf{P} + C \mathbf{I}_K \right]^{-1} (\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K)' \boldsymbol{\eta} + C \right\}, \quad (12)$$

$$\hat{\boldsymbol{\gamma}} = [(\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K)' \mathbf{M} (\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K) + n\lambda \mathbf{P} + \hat{C} \mathbf{I}_K]^{-1} (\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K)' \boldsymbol{\eta}, \quad (13)$$

with $\hat{\boldsymbol{\gamma}}$ initially set to the pilot estimator $\hat{\boldsymbol{\gamma}}_0$. Equation (11) corresponds to the generalized least squares problem that updates $\hat{\boldsymbol{\beta}}$ for a given $\hat{\boldsymbol{\gamma}}$. Equations (12)-(13) correspond to the quadratically constrained quadratic program that updates $\hat{\boldsymbol{\gamma}}$ for a given $\hat{\boldsymbol{\beta}}$ using the method of Lagrange multipliers. The Lagrange multiplier \hat{C} is computed by performing the singular value decomposition of $(\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K)' \mathbf{M} (\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K) + n\lambda \mathbf{P}$ and numerically finding the root of a monotone function (see e.g., [Golub and Van Loan, 2013, chap. 6] for details).

2.3 Inference

In this section we discuss the sampling distribution of the HRF estimates and illustrate how to perform inference on model parameters. As before, we omit the voxel index v from notations for conciseness. Recall that for a given voxel, the HRF shape parameter $\boldsymbol{\gamma}$ is well defined only if at least one condition induces an activation in the voxel. Under this assumption, for a sufficiently large sample size n and sufficiently small smoothing parameter λ , the sampling distributions of $\hat{\boldsymbol{\beta}}$ and $\hat{\boldsymbol{\gamma}}$ can be respectively approximated by $N(\boldsymbol{\beta}, [(\mathbf{I}_L \otimes \boldsymbol{\gamma})' \mathbf{M} (\mathbf{I}_L \otimes \boldsymbol{\gamma})]^{-1})$ and $N(\boldsymbol{\gamma}, [(\boldsymbol{\beta} \otimes \mathbf{I}_K)' \mathbf{M} (\boldsymbol{\beta} \otimes \mathbf{I}_K)]^{-1})$, where the matrices $\hat{\mathbf{V}}_j$ are replaced

by the true covariances \mathbf{V}_j in \mathbf{M} and where $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Further, neglecting the uncertainty about $\boldsymbol{\gamma}$ (and about the covariance matrices \mathbf{V}_j) in the estimation of $\boldsymbol{\beta}$ and vice-versa, we obtain

$$\begin{aligned}\hat{\boldsymbol{\beta}} &\approx N\left(\boldsymbol{\beta}, [(\mathbf{I}_L \otimes \hat{\boldsymbol{\gamma}})' \mathbf{M} (\mathbf{I}_L \otimes \hat{\boldsymbol{\gamma}})]^{-1}\right), \\ \hat{\boldsymbol{\gamma}} &\approx N\left(\boldsymbol{\gamma}, [(\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K)' \mathbf{M} (\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K)]^{-1}\right).\end{aligned}\tag{14}$$

The first result allows us to perform inference to detect activation, as is commonly performed in the GLM setting. The second results allows us to test hypotheses regarding the shape of the HRF. For example, after computing the values of $\boldsymbol{\gamma}$ that correspond to the canonical HRF, we could test for deviations in the shape across the brain based on the fact that $(\hat{\boldsymbol{\gamma}} - \boldsymbol{\gamma})' [(\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K)' \mathbf{M} (\hat{\boldsymbol{\beta}} \otimes \mathbf{I}_K)]^{-1} (\hat{\boldsymbol{\gamma}} - \boldsymbol{\gamma})$ follows a χ^2 distribution with K degrees of freedom. A theoretical justification of (14) is provided in Appendix B.

Recall that the above inference procedures rely on the assumption that the voxel under study is activated in at least one condition (otherwise $\boldsymbol{\gamma}$ is not well defined). This assumption amounts to the fact that at least one of the unconstrained HRF coefficients γ_{lk} is not zero. It can easily be tested using the pilot estimator $\hat{\mathbf{h}}$ of the γ_{lk} defined in (6). More precisely, $\hat{\mathbf{h}}$ has a normal distribution with asymptotic bias zero and asymptotic variance $(\sum_{j=1}^n \mathbf{X}_j' \mathbf{R}_j \mathbf{X}_j)^{-1} (\sum_{j=1}^n \mathbf{X}_j' \mathbf{R}_j \mathbf{V}_j \mathbf{R}_j \mathbf{X}_j) (\sum_{j=1}^n \mathbf{X}_j' \mathbf{R}_j \mathbf{X}_j)^{-1}$ for large n and small λ_0 . If the null hypothesis $\gamma_{lk} = 0$ for all k, l fails to be rejected, then no further inference should be performed for this voxel. Another sensible approach is to use (14) first to test for activations in all voxels and then to infer the HRF shape in activated voxels.

2.4 Simulations

The basic framework for this simulation study is similar to the one found in Lindquist et al. (2008), where a number of different HRF modeling approaches were evaluated. Inside a static brain slice, with dimensions 51×40 , a set of 25 identically sized squares, with dimensions 4×4 , were placed to represent active regions (see Fig. 1A). Within each square, we simulated BOLD fMRI signal based on different stimulus functions, which varied systematically across the squares in terms of onset and duration. From left to right the onset of activation varied from the first to the fifth TR. From top to bottom, the duration of activation varied from one to nine TRs in steps of two. To create the response, we convolved the stimulus function in each square with SPMs canonical HRF, using a modified nonlinear convolution that includes an exponential decay to account for refractory effects with stimulation across time, with the net result that the BOLD response saturates with sustained activity in a manner consistent with observed BOLD responses (Wager et al., 2005). This procedure gave rise to a set of 25 distinct HRF shapes. Fig. 1B shows examples of the 5 HRFs with no onset shift, which are representative of the remaining HRFs.

In total we performed five simulation studies in order to evaluate the properties of the proposed model. Below follows a description of each.

Simulation 1: In this simulation the TR was assumed to be 1s long and the inter-stimulus interval was set to 30s. This activation pattern was repeated to simulate a total of 10 epochs. To simulate a sample of subjects and group random-effects, we generated 15 subject datasets, which consisted of the BOLD time series at each voxel plus white noise, creating a plausible effect size (Cohen’s $d=0.5$) based on observed effect sizes in the visual and motor cortex (Wager et al., 2005). The value of σ_ε^2 was set to 3. In addition, a random between-subject variation with a standard deviation of size one third of the within-subject variation was added to each subject’s time course.

Simulation 2: The second simulated data set was constructed in precisely the same manner as outlined in Simulation 1, except here instead of using white noise to simulate within-subject variation we used an AR(1) noise process with $\theta_\varepsilon = 0.3$.

Simulation 3: The third simulated data set was constructed, using the exact same process with AR(1) noise described in Simulation 2, except here instead of using SPMs canonical HRF we used a subject-specific HRF. These were randomly generated using 20 B-spline basis sets with weights drawn from a normal distribution with mean equal to the weights corresponding to the canonical HRF and standard deviation 0.1.

Simulation 4: The fourth simulated data set was constructed in precisely the same manner as outlined in Simulation 1, except here we allowed for two separate conditions. For both conditions the inter-stimulus interval was set to 30s and the activation pattern was repeated to simulate a total of 10 epochs. However, the two conditions were interleaved to begin 15s apart from one another. The β value for the two conditions were set to 0.5 and 1, respectively. All other parameters were set according to the description of Simulation 1.

Simulation 5: The fifth simulated data set was constructed in precisely the same manner as outlined in Simulation 1, except here we used a fast event-related design. The inter-stimulus interval was randomized across each trial using a uniform distribution between 10 and 20s. All other parameters were set according to the description of Simulation 1.

For each of the five simulations, the basic data sets of dimensions $51 \times 40 \times 300$ were fit voxel-wise using both using the standard GLM/OLS approach and our proposed hierarchical approach. For Simulations 1-3 an event-related stimulus function with a single spike repeated every 30s was used for fitting the models to the data set described above. For Simulation 4 this was supplemented by a second stimulus function with a single spike repeated every 30s corresponding to the timing of the second condition. Finally, in Simulation 5 we used a stimulus function with a single spike repeated according to the outcomes of the randomization scheme outlined above.

This implies that for each simulation the square in the upper left-hand corner of Fig. 1A is correctly specified for the standard GLM while the remaining squares have activation profiles that are mis-specified to various degrees. When fitting our method we used 20 B-spline basis functions of order 6. While our proposed method provides multi-subject

estimates and a framework for direct inference on these parameters, the standard OLS (ordinary least squares) approach to multi-subject analysis in fMRI involves using a two-stage model. It begins by fitting individual regression coefficients for each subject using a standard GLM. Thereafter group estimates of the parameters are obtained by averaging across subjects and variance components are obtained by computing the variance of the estimates across subjects. Since this analysis is performed on the estimated regression coefficients, the variance will contain contributions from both the standard error of the estimates and the between-subject variance components. This method is the most popular method in the neuroimaging community for estimating the parameters of a mixed-effects model [Mumford and Nichols, 2009].

After estimation we performed population level inference to determine whether β was significantly different from 0. In each case we performed a one-sided test. In order to control for multiple comparisons we used an FDR-controlling procedure with $q = 0.05$ [Genovese et al., 2002].

2.5 Experimental Data

All subjects ($n = 20$) provided informed consent in accord with the Declaration of Helsinki, and the Columbia University Institutional Review Board approved all procedures. Subjects were all right-handed, as assessed by the Edinburgh Handedness Inventory, and free of self-reported history of psychiatric and neurological disorders, excessive caffeine or nicotine use, and illicit drug use. They were pre-screened during an initial calibration session to ensure that stimuli were painful and that they could rate pain reliably ($r \geq .65$) between applied temperature and pain rating). During fMRI scanning, 48 thermal stimuli, 12 at each of 4 temperatures, were delivered to the left forearm using a Peltier device (TSA-II, Medoc, Inc.) with an fMRI-compatible 1.5 mm-diameter thermode. Temperatures were calibrated individually for each participant before scanning to be warm, mildly painful, moderately painful, and near tolerance. Heat stimuli were preceded by a 2s warning cue and 6s anticipation period, lasted 10s in duration (1.5s ramp up/down, 7s at peak), and were followed by a 30s inter-trial interval (ITI). At a time 14s into the ITI, participants were asked to rate the painfulness of the stimulus on an 8-point visual analogue scale using an fMRI-compatible trackball (Resonance Technologies, Inc.). Functional T2*-weighted EPI-BOLD images (TR = 2s, $3.5 \times 3.5 \times 4$ mm voxels, ascending interleaved acquisition) were collected during 6 functional runs each consisting of 6 minutes 8s. Images were corrected for slice-timing acquisition delay and realigned to adjust for head motion using SPM5 software (<http://www.fil.ion.ucl.ac.uk/spm/>). A high-resolution anatomical image (T1-weighted spoiled-GRASS [SPGR] sequence, $1 \times 1 \times 1$ mm voxels, TR = 19ms) was collected after the functional runs and coregistered to the mean functional image using a mutual information cost function (SPM5, with manual checking and adjustment of starting values to ensure satisfactory alignment for each participant), and was then segmented and warped to the Montreal Neurologic Institute template (avg152T1.nii) using SPM5’s “generative” segmentation [Ashburner and Friston, 2005]. Warps were applied to functional images. Functional images were smoothed with a 6 mm-FWHM Gaussian kernel, high-pass filtered with a 120s (.0083 Hz) discrete cosine basis set (SPM5), and windsorized at 2 standard deviations prior to analysis. Each of the models described above was fit to data from

each voxel in a single axial slice ($z = -22\text{mm}$) that covered several pain-related regions of interest, including the anterior cingulate cortex. Separate HRFs were estimated for stimuli of each temperature, though we focus on the responses to the highest heat level in the results.

We fit the data using our proposed hierarchical approach with 12 B-spline basis functions of order 6. As the fMRI data consisted of 6 functional runs for each subject, we combined the results across runs in a manner corresponding to a fixed-effects model, assuming the same HRF for all runs for a given subject and voxel. As an alternative it would be possible to extend our model to allow β to be a random effect across runs in a manner analogous to that outlined in [Badillo et al., 2013b]. After estimation we performed population level inference to determine whether β was significantly different from 0. In each case we performed a one-sided test. In order to control for multiple comparisons we used an FDR-controlling procedure with $q = 0.05$ [Genovese et al., 2002].

3 RESULTS

3.1 Simulations

Figs. 2-4 show results of the first three simulation studies. Each contains estimates of β , θ_ε , σ_ε^2 , the t -map, and the thresholded t -map, obtained using both the standard GLM/OLS approach and our hierarchical model. In each case, GLM/OLS (first row) gives reasonable results for delayed onsets within 3s and widths up to 3s, corresponding to squares in the upper left-hand corner. However, its performance worsens dramatically as onset and duration increase. As pointed out in Lindquist et al. (2008) this is natural as the GLM is correctly specified for the square in the upper left-hand corner, but not well equipped to handle large deviations from this model. Interestingly, in Fig. 2 we see that this model misspecification gives rise to a positive autocorrelation in several of the most severely effected squares, even though we used white noise in that particular simulation.

Almost all of these problems are solved using our hierarchical model. Clearly, irrespective of the shape of the underlying HRF we were able to efficiently recover both β and the variance component in each square. Clearly, these improved estimations lead to significantly improved sensitivity and specificity in the population-level hypothesis tests.

The results of Simulation 4 are shown in Fig. 5. Here the first column contains the t -map corresponding to the test of whether the contrast, $\beta_1 - \beta_2$, between the parameters of the two conditions was significantly different from zero, while the second column shows the analogous thresholded t -maps. Again, our hierarchical model is able to detect significant regions with both high sensitivity and specificity. In contrast, the GLM/OLS approach performs poorly, even in the upper-left hand corner where it is expected to be optimal.

Fig. 6 shows results for Simulation 5 where we used a rapid event-related task. The first column contains the t -map corresponding to testing whether the parameter β was significantly different from zero, and the second column corresponds to the thresholded t -

maps. The GLM/OLS approach performs somewhat better than in the other simulations, perhaps due to the increased number of trials. However, our hierarchical model is again able to detect significant regions with both a high degree of sensitivity and specificity.

To illustrate further, we see in Fig. 7 the results of 100 replications of the first simulation. Fig. 7A shows the portion of times the standard GLM/OLS approach gave significant results in each voxel across the simulated brain. Fig. 7B shows the same results for our hierarchical model. Clearly, our hierarchical model is able to consistently detect truly activated regions, while avoiding spurious findings. Not surprisingly, the GLM/OLS approach only performs well in squares where it is correctly defined. Fig 7C shows the portion of times the estimated HRF significantly deviated from SPMs canonical form ($\alpha < 0.05$). This test was performed by first computing values of γ corresponding to the canonical HRF and thereafter using the χ^2 test described in Section 2.3. The results show, again not surprisingly, that the HRF in voxels in the lower right-hand corner of the brain deviate significantly from the canonical form. Interestingly, these voxels corresponds closely to voxels that were erroneously deemed non-active in Fig. 7A. This leads us to believe that our approach could have an alternative use as a diagnostic tool for assessing the performance of standard GLM analyses.

Fig 8 shows the methods ability to recover the time-to-peak and width of the underlying HRF used to generate the data. The left hand column shows the true values and the right hand columns the mean estimated values across the 100 replications. Clearly we are able to extremely accurately capture the true value of the time-to-peak. However, it appears that the estimates of the width, while close are somewhat confounded by the changes in onset, with the best results occurring when there are no onset shifts present.

3.2 Experimental Data

The results of the pain experiment are shown in Fig. 9 for a single axial slice ($z = -22\text{mm}$). The location of the slice used and an illustration of key areas of interest are shown in Fig. 9A. The highlighted areas are the rostral dorsal anterior cingulate cortex (rdACC) and the secondary somatosensory cortex (S2); two brain regions known from previous work to be involved in the processing of pain intensity [Ferretti et al., 2003, Peyron et al., 2000]. Activation in S2 is thought to be related to sensory-discriminant aspects of pain-processing, while rdACC has been shown to be related to expectancy [Atlas et al., 2012].

In Fig. 10 we see examples of the estimated HRFs, obtained using our approach, from voxels chosen because they lay in the center of the rdACC and S2, respectively. These results include both the subject-specific and group-level estimates. The shapes of both group-level HRFs are significantly different from the canonical HRF ($p < 0.01$) using the χ^2 test. Note that the shapes of the HRFs are also quite different from one another. For rdACC, the width of the response is significantly wider from what we would expect from the canonical HRF. While for S2, the onset is significantly delayed. However, interestingly both HRFs appear to reach their peak at roughly the same time point following activation.

Due to the apparent variability in the HRF across the slice, it would be problematic to analyze this data set using a canonical HRF or one that uses a constrained basis set. In

fact, the GLM showed no activation in either rdACC or S2. However, the flexibility of our approach allows for large deviations in the shape of the HRF across voxels. In Fig. 9B we see an activation map obtained using our method. In particular, note that there is significant activation in S2 in response to the noxious stimulation. In previous analysis of this data set [Lindquist et al., 2009], activations in this region were particularly difficult to detect using standard GLM methods (e.g., the canonical HRF plus its temporal and dispersion derivatives, or the finite impulse response (FIR) basis set). The inverse-logit (IL) model was the only approach that showed activation in S2 contralateral to noxious stimulation. The activation shown here is extremely robust in comparison.

4 DISCUSSION

In this paper we introduce a new approach towards the simultaneous detection of activation and the estimation of the HRF for multi-subject fMRI. The suggested approach circumvents a number of shortcomings in the standard approach for performing group analysis. In these approaches there is often a tension between flexible modeling of the HRF at the first level and straightforward inference in the second level. For example, if multiple basis sets are used in the first-level GLM, then it is difficult to determine an appropriate contrast to bring forward to the second-level. Often, researchers use only a subset of the basis functions and therefore potentially ignore important information contained in those left behind. The proposed approach allows the shape of the hemodynamic response function to vary across regions and subjects, as when using basis sets, while still providing a straightforward way to estimate population-level activation using all the information from the first level analysis.

An additional benefit of our model is that the suggested inferential framework not only provides a means for performing the standard tests for determining whether a voxel is significantly active, but also allows one to test whether the estimated HRF deviates from some canonical shape (e.g., SPMs canonical HRF). This type of inference has been under-utilized in the field, but we feel it can be extremely useful for diagnostic purposes, as it can help identify regions that would normally not be deemed active when using a canonical HRF.

In addition, our method could prove useful in situations when the standard HRF is either ill-fitting, such as in studies of young or elderly populations, or when the exact onset time or width of activation is unknown and added flexibility is needed to properly fit the data. An example of the latter is the thermal pain data presented in this paper. In previous studies [Lindquist et al., 2009], activations in S2 were particularly difficult to detect using either the canonical HRF plus its temporal and dispersion derivatives, or the finite impulse response (FIR) basis set. However, using our approach, the activation was extremely robust in comparison.

To date the method has only been implemented in a voxel-size manner. Hence, we make the common, but rather implausible, assumption of independence between voxels with regards to both the HRF shape and amplitude. However, we do allow for the possibility of a parcel-specific noise structure. We are currently in the process of extending the method to also

estimate spatial dependences in the data. For that purpose two types of approaches can be considered: one consists in spatially regularizing the estimation in local neighborhoods; the other is to integrate a functional parcellation of the brain. In the latter case, one can either resort to an existing atlas (see e.g. [Karahanoğlu et al., 2013]) or define a data-driven parcellation (see e.g., [Chaari et al., 2012] for such a parcellation at the subject level). Though to make the manuscript more manageable we decided to limit our discussion of this issue until a later time.

The presented simulations and data set are identical to those used in a previous study [Lindquist et al., 2009], where we evaluated the performance of seven different HRF models. These included: SPMs canonical HRF; the canonical HRF plus its temporal derivative; the canonical HRF plus its temporal and dispersion derivatives; the finite impulse response (FIR) basis set; a regularized version of the FIR model (denoted the smooth FIR); a nonlinear model with the same functional form as the canonical HRF but with 6 variable parameters; and the inverse logit (IL) model. The results of that work showed that it was surprisingly difficult to accurately recover true task-evoked changes in BOLD signal and that there were substantial differences among models in terms of power, bias and parameter confusability. While the derivative models were accurate for very short shifts in latency they became progressively less accurate as the shift increased. The IL model and the smooth FIR model showed the least amount of biases, and the IL model showed by far the least amount of confusability of all the models that were examined. Both these methods were clearly able to handle even large amounts of model misspecification and uncertainty about the exact timing of the onset and duration of activation.

The suggested model clearly outperformed each of the 7 other models in the same battery of tests. For space purposes we only present the canonical HRF for comparison purposes, and we encourage interested readers to look back at [Lindquist et al., 2009], for more results. The hierarchical model was found to have a superb balance of sensitivity and specificity that none of the other models was able to obtain. In addition, in the previous work the IL model was the only model that showed significant activation in S2 contralateral to noxious stimulation. Here the proposed model shows extremely robust signal in this region. For these reasons we believe the proposed model is a useful approach towards effectively modeling multi-subject fMRI data.

A Matlab implementation of the proposed methodology is available upon request. The code runs in a reasonable time for fMRI data sets of moderate size. By optimizing the code and running it in parallel on voxels and/or on subjects, our methodology can scale up to large fMRI data. Specifically, step 1 of the estimation algorithm can be solved in closed form and requires few matrix multiplications. Step 2 is also very fast due to the definition of the noise parameters at the parcel level and to the computational efficiency of the Yule-Walker equations. Step 3 (maximum likelihood/EM algorithm) would be very slow if it was run for each voxel but is in reality only carried for a small number of voxels. Step 4 consists in a large number of standard quadratic programming problems that can be solved quickly and in parallel. Step 5 is arguably the slowest part of the estimation procedure because of the necessity to compute inverse data covariance matrices. However, linear algebra tricks can reduce the dimension of the matrices to be inverted from the number of scans (T_j) to the number of regressors (KL).

Acknowledgements

We thank two anonymous reviewers for remarks that helped us improve the quality of this paper, and Tor Wager for supplying the data. This research was partially supported by NIH grant R01EB016061.

References

- [Aguirre et al., 1998] Aguirre, G. K., Zarahn, E., and D’Esposito, M. (1998). The variability of human, BOLD hemodynamic responses. *NeuroImage*, 8(4):360–369.
- [Ashburner and Friston, 2005] Ashburner, J. and Friston, K. J. (2005). Unified segmentation. *Neuroimage*, 26(3):839–851.
- [Atlas et al., 2012] Atlas, L. Y., Whittington, R. A., Lindquist, M. A., Wielgosz, J., Sonty, N., and Wager, T. D. (2012). Dissociable influences of opiates and expectations on pain. *The Journal of Neuroscience*, 32(23):8053–8064.
- [Badillo et al., 2013a] Badillo, S., Vincent, T., and Ciuciu, P. (2013a). Group-level impacts of within- and between-subject hemodynamic variability in fMRI. *NeuroImage*, 82(0):433 – 448.
- [Badillo et al., 2013b] Badillo, S., Vincent, T., and Ciuciu, P. (2013b). Multi-session extension of the joint-detection framework in fMRI. In *ISBI*, pages 1512–1515.
- [Brockwell and Davis, 2006] Brockwell, P. J. and Davis, R. A. (2006). *Time series: theory and methods*. Springer Series in Statistics. Springer, New York. Reprint of the second (1991) edition.
- [Calhoun et al., 2004] Calhoun, V. D., Stevens, M. C., Pearlson, G. D., and Kiehl, K. A. (2004). fMRI analysis with the general linear model: removal of latency-induced amplitude bias by incorporation of hemodynamic derivative terms. *NeuroImage*, 22:252–257.
- [Chaari et al., 2012] Chaari, L., Forbes, F., Vincent, T., and Ciuciu, P. (2012). Hemodynamic-informed parcellation of fmri data in a joint detection estimation framework. In Ayache, N., Delingette, H., Golland, P., and Mori, K., editors, *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2012*, volume 7512 of *Lecture Notes in Computer Science*, pages 180–188. Springer Berlin Heidelberg.
- [Chaari et al., 2013] Chaari, L., Vincent, T., Forbes, F., Dojat, M., and Ciuciu, P. (2013). Fast joint detection-estimation of evoked brain activity in event-related fMRI using a variational approach. *IEEE Transactions on Medical Imaging*, 32(5):821–837.
- [Demidenko, 2004] Demidenko, E. (2004). *Mixed models*. Wiley Series in Probability and Statistics. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ. Theory and applications.
- [Ferretti et al., 2003] Ferretti, A., Babiloni, C., Gratta, C. D., Caulo, M., Tartaro, A., Bonomo, L., Rossini, P. M., and Romani, G. L. (2003). Functional topography of the

- secondary somatosensory cortex for nonpainful and painful stimuli: an fMRI study. *Neuroimage*, 20(3):1625–1638.
- [Friston et al., 1998] Friston, K., Fletcher, P., Josephs, O., Holmes, A., Rugg, M., and Turner, R. (1998). Event-related fMRI: characterizing differential responses. *Neuroimage*, 7(1):30–40.
- [Genovese et al., 2002] Genovese, C., Lazar, N., and Nichols, T. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *NeuroImage*, 15:870–878.
- [Genovese, 2000] Genovese, C. R. (2000). A bayesian time-course model for functional magnetic resonance imaging data. *Journal of the American Statistical Association*, 95(451):691–703.
- [Glover, 1999] Glover, G. H. (1999). Deconvolution of impulse response in event-related BOLD fMRI. *NeuroImage*, 9(4):416–429.
- [Golub and Van Loan, 2013] Golub, G. H. and Van Loan, C. F. (2013). *Matrix computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, fourth edition.
- [Goutte et al., 2000] Goutte, C., Nielsen, F. A., and Hansen, L. K. (2000). Modeling the haemodynamic response in fMRI using smooth FIR filters. *IEEE Transactions on Medical Imaging*, 19:1188–1201.
- [Handwerker et al., 2004] Handwerker, D. A., Ollinger, J. M., and D’Esposito, M. (2004). Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *NeuroImage*, 21(4):1639 – 1651.
- [Karahanoğlu et al., 2013] Karahanoğlu, F. I., Caballero-Gaudes, C., Lazeyras, F., and Ville, D. V. D. (2013). Total activation: fmri deconvolution through spatio-temporal regularization. *NeuroImage*, 73(0):121 – 134.
- [Kiehl and Liddle, 2001] Kiehl, K. A. and Liddle, P. F. (2001). An event-related functional magnetic resonance imaging study of an auditory oddball task in schizophrenia. *Schizophrenia Research*, 48(2):159–171.
- [Liao et al., 2002] Liao, C., Worsley, K. J., Poline, J.-B., Duncan, G. H., and Evans, A. C. (2002). Estimating the delay of the response in fMRI data. *NeuroImage*, 16:593–606.
- [Lindquist et al., 2009] Lindquist, M., Loh, J., Atlas, L., and Wager, T. (2009). Modeling the hemodynamic response function in fMRI: Efficiency, bias and mis-modeling. *NeuroImage*, 45(1, Supplement 1):S187 – S198.
- [Lindquist, 2008] Lindquist, M. A. (2008). The statistical analysis of fMRI data. *Statistical Science*, 23(4):439–464.
- [Lindquist and Wager, 2007] Lindquist, M. A. and Wager, T. D. (2007). Validity and power in hemodynamic response modeling: A comparison study and a new approach. *Human Brain Mapping*, 28:764–784.

- [Makni et al., 2005] Makni, S., Ciuciu, P., Idier, J., and Poline, J.-B. (2005). Joint detection-estimation of brain activity in functional MRI: a multichannel deconvolution solution. *IEEE Transactions on Signal Processing*, 53(9):3488–3502.
- [Makni et al., 2008] Makni, S., Idier, J., Vincent, T., Thirion, B., Dehaene-Lambertz, G., and Ciuciu, P. (2008). A fully bayesian approach to the parcel-based detection-estimation of brain activity in fMRI. *Neuroimage*, 41(3):941–969.
- [Mumford and Nichols, 2009] Mumford, J. A. and Nichols, T. (2009). Simple group fMRI modeling and inference. *Neuroimage*, 47(4):1469–1475.
- [Nocedal and Wright, 2006] Nocedal, J. and Wright, S. J. (2006). *Numerical optimization*. Springer Series in Operations Research and Financial Engineering. Springer, New York, second edition.
- [Pawitan, 2001] Pawitan, Y. (2001). *In all likelihood: statistical modelling and inference using likelihood*. Oxford science publications. Clarendon press, Oxford.
- [Peyron et al., 2000] Peyron, R., Laurent, B., and Garcia-Larrea, L. (2000). Functional imaging of brain responses to pain. A review and meta-analysis (2000). *Neurophysiologie Clinique/Clinical Neurophysiology*, 30(5):263–288.
- [Poldrack et al., 2011] Poldrack, R. A., Mumford, J. A., and Nichols, T. E. (2011). *Handbook of functional MRI data analysis*. Cambridge University Press.
- [Riera et al., 2004] Riera, J. J., Watanabe, J., Kazuki, I., Naoki, M., Aubert, E., Ozaki, T., and Kawashima, R. (2004). A state-space model of the hemodynamic approach: nonlinear filtering of bold signals. *NeuroImage*, 21:547–567.
- [Sanyal and Ferreira, 2012] Sanyal, N. and Ferreira, M. A. (2012). Bayesian hierarchical multi-subject multiscale analysis of functional mri data. *Neuroimage*, 63(3):1519–1531.
- [Schacter et al., 1997] Schacter, D. L., Buckner, R. L., Koutstaal, W., Dale, A. M., and Rosen, B. R. (1997). Rectangular confidence regions for the means of multivariate normal distributions. *Late onset of anterior prefrontal activity during true and false recognition: an event-related fMRI study*, 6:259–269.
- [Vincent et al., 2010] Vincent, T., Risser, L., and Ciuciu, P. (2010). Spatially adaptive mixture modeling for analysis of fmri time series. *Medical Imaging, IEEE Transactions on*, 29(4):1059–1074.
- [Woolrich et al., 2004] Woolrich, M. W., Behrens, T. E., and Smith, S. M. (2004). Constrained linear basis sets for HRF modelling using variational Bayes. *NeuroImage*, 21(4):1748–1761.
- [Worsley and Friston, 1995] Worsley, K. J. and Friston, K. J. (1995). Analysis of fMRI time-series revisited-again. *NeuroImage*, 2:173–181.
- [Zarahn, 2002] Zarahn, E. (2002). Using larger dimensional signal subspaces to increase sensitivity in fMRI time series analyses. *Human brain mapping*, 17(1):13–16.

[Zhang et al., 2013] Zhang, T., Li, F., Beckes, L., and Coan, J. A. (2013). A semi-parametric model of the hemodynamic response for multi-subject fmri data. *NeuroImage*, 75(0):136 – 145.

Appendix

A Estimation of the dependence in subject effects

A.1 EM algorithm

We fix a voxel and omit the index v from notations. Since the random effects $\boldsymbol{\xi}_j$ are independent, identically distributed, and stationary, it holds that $\sigma_{\xi l}^2 = \mathbb{E}(\|\boldsymbol{\xi}_{jl}\|^2)/K$ for each subject ($1 \leq j \leq n$) and condition ($1 \leq l \leq L$). Hence, we first estimate $\sigma_{\xi l}^2$ by the empirical average $\sigma_{\xi l}^{2(0)} = (1/n) \sum_{j=1}^n \|\hat{\boldsymbol{\xi}}_{jl}\|^2/K$. We initially assume working independence for the $\boldsymbol{\xi}_{jl}$ so that $\rho_{\xi l}^{(0)}(k) = \delta_{0k}$ for each lag ($1 \leq k \leq K-1$) and condition. At the $(r+1)^{th}$ iteration ($r \geq 0$) of the EM algorithm, the E-step computes the conditional expectation of (minus twice the logarithm of) the complete likelihood, i.e., the likelihood of the augmented data $(\mathbf{y}_1, \dots, \mathbf{y}_n, \boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_n)$. The conditioning variables are $\mathbf{y}_1, \dots, \mathbf{y}_n$, the current estimators $\boldsymbol{\sigma}_{\xi}^{2(r)}, \boldsymbol{\rho}_{\xi}^{(r)}$, and $\hat{\beta}_0, \hat{\gamma}_0, \hat{\sigma}_{\varepsilon m}^2, \hat{\theta}_{\varepsilon m}$. Up to constant terms, the conditional expectation is

$$Q(\boldsymbol{\sigma}_{\xi}^2, \boldsymbol{\rho}_{\xi} | \boldsymbol{\sigma}_{\xi}^{2(r)}, \boldsymbol{\rho}_{\xi}^{(r)}) = nK \ln |\mathbf{D}_{\xi}| + n \ln |\mathbf{T}_{\xi}| \\ + \sum_{j=1}^n \boldsymbol{\xi}_j^{(r)'} \mathbf{T}_{\xi}^{-1} (\mathbf{D}_{\xi}^{-1} \otimes \mathbf{I}_K) \boldsymbol{\xi}_j^{(r)} + \sum_{j=1}^n \text{tr}\{\mathbf{T}_{\xi}^{-1} (\mathbf{D}_{\xi}^{-1} \otimes \mathbf{I}_K) \mathbf{B}_j^{(r)}\}, \quad (15)$$

where $\mathbf{B}_j^{(r)} = [\mathbf{X}_j' \hat{\mathbf{V}}_{\varepsilon jm}^{-1} \mathbf{X}_j + (\mathbf{T}_{\xi}^{(r)})^{-1} ((\mathbf{D}_{\xi}^{(r)})^{-1} \otimes \mathbf{I}_K)]^{-1}$, $\boldsymbol{\xi}_j^{(r)} = \mathbf{B}_j^{(r)} \mathbf{X}_j' \hat{\mathbf{V}}_{\varepsilon jm}^{-1} \mathbf{r}_j$ is the predicted random effect for the j^{th} subject, and \mathbf{r}_j is the residual vector defined in step 2 of section 2.2.

For $1 \leq l \leq L$, the derivative of Q with respect to $\sigma_{\xi l}^2$ is

$$\frac{\partial Q}{\partial \sigma_{\xi l}^2} = \frac{nK}{\sigma_{\xi l}^2} - \frac{1}{\sigma_{\xi l}^4} \sum_{j=1}^n \boldsymbol{\xi}_{jl}^{(r)'} \mathbf{T}_{\xi l}^{-1} \boldsymbol{\xi}_{jl}^{(r)} - \frac{1}{\sigma_{\xi l}^4} \sum_{j=1}^n \text{tr}\{\mathbf{T}_{\xi l}^{-1} \mathbf{B}_{jl}^{(r)}\}, \quad (16)$$

where the matrix $\mathbf{B}_j^{(r)}$ has been partitioned in blocks $\mathbf{B}_{jll'}^{(r)}$, $1 \leq l, l' \leq L$, of size $K \times K$.

Writing $\mathbf{C}_l^{(r)} = \sum_{j=1}^n (\boldsymbol{\xi}_{jl}^{(r)} \boldsymbol{\xi}_{jl}^{(r)'} + \mathbf{B}_{jll}^{(r)})$ and equating (16) with zero, we get

$$\sigma_{\xi l}^2 = \frac{1}{nK} \left(\sum_{j=1}^n \boldsymbol{\xi}_{jl}^{(r)'} \mathbf{T}_{\xi l}^{-1} \boldsymbol{\xi}_{jl}^{(r)} + \sum_{j=1}^n \text{tr}\{\mathbf{T}_{\xi l}^{-1} \mathbf{B}_{jll}^{(r)}\} \right) \\ = \frac{1}{nK} \text{tr}(\mathbf{T}_{\xi l}^{-1} \mathbf{C}_l^{(r)}). \quad (17)$$

After plugging (17) in (15), the variance-profile function Q_p to be optimized in the M step

of the algorithm is

$$\begin{aligned} Q_p(\boldsymbol{\rho}_\xi | \boldsymbol{\sigma}_\xi^{2(r)}, \boldsymbol{\rho}_\xi^{(r)}) &= nK \ln |\mathbf{D}_\xi| + n \ln |\mathbf{T}_\xi| + n \\ &= nK \sum_{l=1}^L \ln (\text{tr}(\mathbf{T}_{\xi l}^{-1} \mathbf{C}_l^{(r)})) - nK \ln(nK) + n \sum_{l=1}^L \ln |\mathbf{T}_{\xi l}| + n. \end{aligned} \quad (18)$$

For computational speed, we may use any suitable gradient-based optimization method. Let \mathbf{D}_k be the $K \times K$ matrix whose (i, j) entry is 1 if $|i - j| = k$ and 0 otherwise. Then $\mathbf{T}_{\xi l} = \sum_{k=0}^{K-1} \rho_{\xi l}(k) \mathbf{D}_k$ and the gradient of Q_p is

$$\frac{\partial Q_p}{\partial \rho_{\xi l}(k)} = -nK \frac{\text{tr}(\mathbf{T}_{\xi l}^{-1} \mathbf{D}_k \mathbf{T}_{\xi l}^{-1} \mathbf{C}_l^{(r)})}{\text{tr}(\mathbf{T}_{\xi l}^{-1} \mathbf{C}_l^{(r)})} + n \text{tr}(\mathbf{T}_{\xi l}^{-1} \mathbf{D}_k) \quad (19)$$

for $1 \leq k \leq K-1$ and $1 \leq l \leq L$. The box constraints $-1 \leq \rho_{\xi l}(k) \leq 1$ and positive definiteness of $\mathbf{T}_{\xi l}$ are enforced during the optimization. The updated variance estimator $\sigma_{\xi l}^{2(r+1)}$ is obtained by plugging $\rho_{\xi}^{(r+1)}$ in (17). Writing $\mathbf{\Gamma}^{(r)} = (\text{vec}(\mathbf{T}_{\xi 1}^{-1} \mathbf{C}_1^{(r)} \mathbf{T}_{\xi 1}^{-1}), \dots, \text{vec}(\mathbf{T}_{\xi L}^{-1} \mathbf{C}_L^{(r)} \mathbf{T}_{\xi L}^{-1}))$, $\mathbf{D} = (\text{vec}(\mathbf{D}_1), \dots, \text{vec}(\mathbf{D}_{K-1}))$, $\mathbf{t}^{(r)} = (\text{tr}(\mathbf{T}_{\xi 1}^{-1} \mathbf{C}_1^{(r)}), \dots, \text{tr}(\mathbf{T}_{\xi L}^{-1} \mathbf{C}_L^{(r)}))'$, and $\mathbf{S} = (\text{vec}(\mathbf{T}_{\xi 1}^{-1}), \dots, \text{vec}(\mathbf{T}_{\xi L}^{-1}))$, the gradient (19) can be compactly written as

$$\frac{\partial Q_p}{\partial \boldsymbol{\rho}_\xi} = -nK \frac{\text{vec}(\mathbf{D}' \mathbf{\Gamma}^{(r)})}{\mathbf{t}^{(r)} \otimes \mathbf{1}_{K-1}} + n \text{vec}(\mathbf{D}' \mathbf{S}),$$

where the division is taken element-wise and $\mathbf{1}_{K-1}$ is a vector containing $(K-1)$ ones.

A.2 Maximum Likelihood Estimation

For a given voxel v , the partial derivatives of the likelihood function (7) are

$$\frac{\partial L}{\partial \sigma_{\xi l}^2} = \sum_{j=1}^n \text{tr}(\mathbf{X}_{jl}' \mathbf{V}_j^{-1}(v) \mathbf{X}_{jl} \mathbf{T}_{\xi l}) - \sum_{j=1}^n \mathbf{r}_j' \mathbf{V}_j^{-1}(v) \mathbf{X}_{jl} \mathbf{T}_{\xi l} \mathbf{X}_{jl}' \mathbf{V}_j^{-1}(v) \mathbf{r}_j(v) \quad (20)$$

and

$$\frac{\partial L}{\partial \rho_{\xi l}(k)} = \sigma_{\xi l}^2 \sum_{j=1}^n \text{tr}(\mathbf{X}_{jl}' \mathbf{V}_j^{-1}(v) \mathbf{X}_{jl} \mathbf{D}_k) - \sigma_{\xi l}^2 \sum_{j=1}^n \mathbf{r}_j' \mathbf{V}_j^{-1}(v) \mathbf{X}_{jl} \mathbf{D}_k \mathbf{X}_{jl}' \mathbf{V}_j^{-1}(v) \mathbf{r}_j(v) \quad (21)$$

for $1 \leq k \leq K-1$ and $1 \leq l \leq L$. Based on (20)-(21) and a suitable gradient-based optimization procedure, we obtained the ML estimators $\hat{\sigma}_{\xi l}^2(v)$ and $\hat{\rho}_{\xi l}(k, v)$. We then define the aggregated correlation estimate $\hat{\rho}_{\xi l}(k)$ as the median of the $\hat{\rho}_{\xi l}(k, v)$ across voxels where the estimation was carried.

B Sampling distribution of the HRF estimators

Here we provide the theoretical justification of the large-sample approximation (14) to the sampling distribution of the estimators $\hat{\beta}(v)$ and $\hat{\gamma}(v)$.

First, the estimators used in this paper rely on standard statistical procedures whose consistency properties are well documented in the literature. In step 1, the penalized least squares estimator $\hat{\mathbf{h}}(v)$ of the HRF coefficients $\gamma_{lk}(v)$ is consistent as $n \rightarrow \infty$ and $\lambda_0 \rightarrow 0$. Note that increasing the number of scans T_j reduces the influence of the noise ε on the estimation but not the sampling variability (subject effects ξ). In fact, the variance of the pilot estimator $\hat{\mathbf{h}}(v)$ is dominated by the sampling variability: it is of order $\mathcal{O}(\sum_j T_j^2 / (\sum_j T_j)^2)$, i.e. $\mathcal{O}(1/n)$ if the T_j are of comparable size. Also, the parameter λ_0 governing the penalty on HRF shapes lying outside the null space Ψ must go to zero to render the pilot estimator asymptotically unbiased. In step 2, the Yule-Walker estimators of the noise parameters $\sigma_{\varepsilon m}^2$ and $\theta_{\varepsilon m}$ are consistent as $\max_j T_j \rightarrow \infty$. The consistency of steps 3-5 derives from the large-sample properties of least squares- and maximum likelihood estimators and from the consistency of the previous estimation steps. Note that the estimators of the noise and random effects parameters are averaged across subjects and space (except for the variance estimators $\hat{\sigma}_{\xi l}^2(v)$). As a consequence, their variance is generally small in comparison to the variance of the HRF estimators.

We now turn to the large-sample distribution of $\hat{\beta}(v)$ and $\hat{\gamma}(v)$. Assuming that for each subject, stimuli of each type are presented sufficiently often, the design matrices \mathbf{X}_j are of order $\mathcal{O}(\sqrt{T_j})$ in norm and the matrix $\mathbf{M} = \mathbf{M}(v)$ is of order $\mathcal{O}(\sum_j T_j)$ in probability. Given the consistency of $\hat{\gamma}(v)$, $\hat{\mathbf{V}}_j(v)$, the normality of the data, and the fact that $\text{Var}(\boldsymbol{\eta}) \approx \mathbf{M}$, one can apply Slutsky's theorem and the Law of Large Numbers in (11) to obtain the large-sample approximation $\hat{\beta}(v) \rightarrow N(\beta(v), [(\mathbf{I}_L \otimes \boldsymbol{\gamma}(v))' \mathbf{M} (\mathbf{I}_L \otimes \boldsymbol{\gamma}(v))]^{-1})$ (with $\hat{\mathbf{V}}_j^{-1}(v)$ in \mathbf{M} replaced by $\hat{\mathbf{V}}_j^{-1}(v)$) as $n, \sum_j T_j \rightarrow \infty$ and $\lambda \rightarrow 0$. Turning to the estimator $\hat{\gamma}(v)$, simple algebraic manipulations in (11) and (13) show that $\hat{C} = -n\lambda \hat{\gamma}(v)' \mathbf{P} \hat{\gamma}(v)$. In other words, the terms $\hat{C} \mathbf{I}_K$ and $n\lambda \mathbf{P}$ in (13) are of order $\mathcal{O}(n\lambda)$ in probability and thus negligible in comparison to \mathbf{M} . Applying the same arguments as with $\hat{\beta}(v)$ in (11), we obtain the limit distribution $\hat{\gamma}(v) \rightarrow N(\boldsymbol{\gamma}(v), [(\boldsymbol{\beta}(v) \otimes \mathbf{I}_K)' \mathbf{M} (\boldsymbol{\beta}(v) \otimes \mathbf{I}_K)]^{-1})$.




Figure 1: Overview of simulation set-up (A) A set of 25 equally sized squares were placed within a static brain image to represent regions of interest. BOLD signals were simulated based on different stimulus functions, which varied systematically across the squares in their onset and duration of neuronal activation. From left to right the onset of activation varied between the squares from the first to the fifth TR. From top to bottom, the duration of activation varied from one to nine TR in steps of two. (B) The five HRFs with varying duration. The plot illustrates differences in time-to-peak and width attributable to changes in duration.

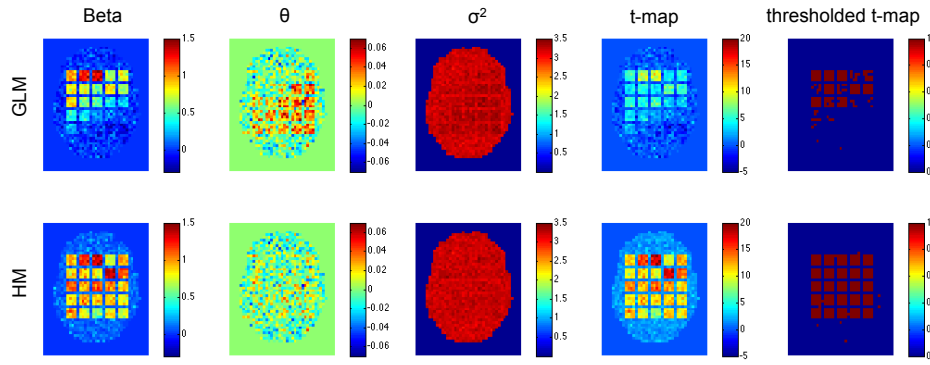


Figure 2: Results of the first simulation shown for the standard GLM/OLS approach (top row), and our hierarchical model (bottom row). From left-to-right the columns represent the estimated values of β , θ_ε , σ_ε^2 , t -map and thresholded t -map.




Figure 3: Results of the second simulation shown for the standard GLM/OLS approach (top row), and our hierarchical model (bottom row). From left-to-right the columns represent the estimated values of β , θ_ε , σ_ε^2 , the t -map and the thresholded t -map.




Figure 4: Results of the third simulation shown for the standard GLM/OLS approach (top row), and our hierarchical model (bottom row). From left-to-right the columns represent the estimated values of β , θ_ε , σ_ε^2 , t -map and thresholded t -map.




Figure 5: Results of the fourth simulation shown for the standard GLM/OLS approach (top row), and our hierarchical model (bottom row). From left-to-right the columns represent the estimated and thresholded t -map for testing whether the contrast between the two conditions was significantly different from 0.




Figure 6: Results of the fifth simulation shown for the standard GLM/OLS approach (top row), and our hierarchical model (bottom row). From left-to-right the columns represent the estimated and thresholded t -map.




Figure 7: The results of 100 replications of the first simulation. (A) The portion of times the standard GLM/OLS approach gave significant results in each voxel. (B) The same results for our hierarchical model. Clearly the hierarchical model is able to effectively separate signal from noise in a more consistent manner than the GLM/OLS. (C) The portion of times the estimated HRF deviated from the canonical form using our model.




Figure 8: The results of 100 replications of the first simulation. (Top row) The true and estimated values of the time-to-peak for the group-level HRF. (Bottom row) Same results for the width.




Figure 9: (A) The location of the slice and an illustration of areas of interest. Both rdACC and S2 are regions known to process pain intensity. (B) A statistical map obtained using the proposed hierarchical model.




Figure 10: Estimates of the subject-specific HRF computed using voxels from the rdACC and S2. The group-level estimates are shown in bold.